

Stop or Continue Data Collection: A Nonignorable Missing Data Approach for Continuous Variables

Thais Paiva¹ and Jerome P. Reiter²

¹Department of Computer Science, Federal University of Minas Gerais, Belo Horizonte, Brazil

²Department of Statistical Science, Duke University, Durham, NC, USA

November 9, 2015

Abstract

We present an approach to inform decisions about nonresponse followup sampling. The basic idea is (i) to create completed samples by imputing nonrespondents' data under various assumptions about the nonresponse mechanisms, (ii) take hypothetical samples of varying sizes from the completed samples, and (iii) compute and compare measures of accuracy and cost for different proposed sample sizes. As part of the methodology, we present a new approach for generating imputations for multivariate continuous data with nonignorable unit nonresponse. We fit mixtures of multivariate normal distributions to the respondents' data, and adjust the probabilities of the mixture components to generate nonrespondents' distributions with desired features. We illustrate the approaches using data from the 2007 U. S. Census of Manufactures.

Keywords: adaptive, missing, mixture, nonignorable, nonresponse

1 Introduction

With shrinking budgets and expanding costs, government agencies and survey organizations must find ways to reduce the expenses of data collection. One promising set of approaches is adaptive design, in which the agency changes the design of the survey based on the characteristics of the respondents and any known information about the nonrespondents. For example, the agency may focus followup efforts on individuals with characteristics not adequately represented among the respondents (e.g., Wagner, 2008; Schouten et al., 2009, 2011, 2012), or cease followup efforts for individuals deemed highly unlikely to respond without excessive expense (e.g., Miller, 2013; Finamore et al., 2013).

In this article, we consider one potential decision point in adaptive designs: should the agency stop all efforts at data collection, given what has been already obtained? Stopping data collection reduces costs and allows for earlier release of data, but doing so only makes

sense when data quality is not sacrificed to undesirable levels. To make informed decisions, agencies can benefit from principled approaches for assessing the trade offs in projected costs and estimated accuracy at specific times in the data collection process.

Several methods for stopping data collection have been suggested in the literature. Most relevant to our work, Rao et al. (2008) and Wagner and Raghunathan (2010) propose stopping rules for surveys with multiple waves of nonresponse followup and a binary response variable. Their basic framework is to impute the missing values for all individuals who have yet to respond by wave t , evaluate some estimand on the completed data, and stop when the change in the estimate from wave $t - 1$ to wave t is small according to some criterion. Rao et al. (2008) and Wagner and Raghunathan (2010) use imputation models that presume the nonrespondents at wave t follow the same distribution as respondents up to (or possibly only in) wave $t - 1$; that is, the nonrespondents' data are missing at random (MAR) (Rubin, 1976). While convenient, MAR is a strong assumption that may be unrealistic, resulting in unreliable imputations and, consequently, ineffective stopping decisions.

In this article, we propose an imputation-based approach to defining stopping rules when the key variables of interest are continuous in nature. Unlike Rao et al. (2008) and Wagner and Raghunathan (2010), we create completed datasets without assuming MAR; that is, we develop and use a multivariate model for nonignorable unit nonresponse. To do so, we fit mixtures of multivariate normal distributions to the respondents' data, and adjust the probabilities of the mixture components (keeping fixed the location and scale parameters) to generate nonrespondents' distributions with desired features. For example, to represent a scenario in which nonrespondents are more (less) likely than respondents to come from a particular region of support, we inflate (deflate) the probabilities of the components with high density in that region. These adjustments facilitate sensitivity analysis—arguably the best one can do with nonignorable missing data—as the agency can evaluate inferences on completed data constructed with different sets of altered probabilities. With such evaluations, agencies can decide whether or not the potential cost savings are worth the loss in accuracy from stopping.

The remainder of this article is organized as follows. In Section 2, we describe how we adapt mixtures of multivariate normal distributions to handle nonignorable unit nonresponse. We note that this approach could be used for general imputation tasks with multivariate continuous data; it is not restricted to stopping rule decisions. In Section 3, we illustrate how to apply this methodology on a subset of data from the 2007 U. S. Census of Manufactures. In Section 4, we describe how to use the nonignorable missing data procedure and sensitivity analysis to inform stopping decisions. We define several accuracy metrics that can aid the decision process. In Section 5, we illustrate the stopping rule methodology on the data from the 2007 U. S. Census of Manufactures. Finally, in Section 6 we conclude with a discussion and suggestions for future research.

2 Model for Multivariate Continuous Data with Non-ignorable Nonresponse

To fix notation, consider an agency that intends to collect data on n individuals, either in a census or a sample. For $i = 1, \dots, n$, let y_{ij} denote the value of variable j for individual i . Let $\mathbf{y}_i = (y_{i1}, \dots, y_{ip})$ include the values of the p variables of interest. Let $\mathbf{Y} = (\mathbf{y}_1, \dots, \mathbf{y}_n)'$. At any time t in the data collection period, let $r_{it} = 1$ if individual i has provided data by time t , and let $r_{it} = 0$ if individual i has not yet responded. Let $\mathbf{R}_t = (r_{1t}, \dots, r_{nt})$. We only consider unit nonresponse when defining each r_{it} . Let \mathbf{Y}_t comprise the data for the n_t individuals with $r_{it} = 1$. We note that $(\mathbf{R}_t, \mathbf{Y}_t)$ is updated as t advances. At any time t , our goal is to impute values of \mathbf{Y} for the nonrespondents, allowing for possibly nonignorable missingness.

When handling multivariate data that are not missing at random (NMAR), analysts generally take one of two approaches. In selection models, analysts posit a joint model for \mathbf{Y} and a model for $(\mathbf{R}_t | \mathbf{Y})$, such as a logistic regression (Glynn et al., 1986; Rubin, 1987; Diggle and Kenward, 1994; Molenberghs et al., 1997). In pattern mixture models, analysts posit a model for \mathbf{R}_t and a model for $(\mathbf{Y} | \mathbf{R}_t)$; that is, there are different models for \mathbf{Y} for respondents and nonrespondents (Glynn et al., 1993; Little, 1993, 1994; Molenberghs et al., 1998; Thijs et al., 2002). We adopt the pattern mixture approach, as we find it easier to utilize for sensitivity analysis; see Little and Rubin (2002), Little (1995, 2008), and Daniels and Hogan (2008) for further discussion.

Pattern mixture models can be implemented in two steps. First, the analyst fits a model to the observed data, say $f(\mathbf{Y}_t)$. Second, the analyst imposes departures from $f(\mathbf{Y}_t)$, for example by altering certain model parameters, that reflect beliefs about possible distributions for the nonrespondents. The analyst can generate multiple imputations (Rubin, 1987) for nonrespondents to create completed versions of \mathbf{Y} for inference. By examining multiple possible distributions for the nonrespondents, the analyst can perform sensitivity analysis.

When building off $f(\mathbf{Y}_t)$, it is beneficial to use a model that accurately describes the multivariate distribution among the respondents. Often this multivariate distribution has complicated features, especially for economic data. We use mixtures of multivariate normal distributions as the engine for flexible modeling of the joint distribution of \mathbf{Y}_t . These models can capture complex relationships among variables, such as non-linearities and interactions, as well as multiple modes, heteroscedasticity, and skewness (Ferguson, 1983; Ghosal and van der Vaart, 2001, 2007; Müller and Mitra, 2013). These models also have parameters that can be easily manipulated to generate possible distributions for nonrespondents.

2.1 Mixtures of multivariate normal distributions

To begin, we review the mixture normal distribution that we use to model \mathbf{Y}_t at any time t . To facilitate model specification and estimation, we first standardize each variable in \mathbf{Y} . We assume that each individual in the dataset belongs to one of $K < \infty$ mixture components, indicated by $z_i \in \{1, \dots, K\}$. Let $\mathbf{Z} = (z_1, \dots, z_n)$. Let $\pi_k = P(z_i = k)$, where

$i = 1, \dots, n$ and $k = 1, \dots, K$, be the probability that individual i belongs to component k . Let $\boldsymbol{\pi} = (\pi_1, \dots, \pi_K)$. Within each component, let \mathbf{y}_i follow a multivariate normal distribution with mean $\boldsymbol{\mu}_k$ and variance $\boldsymbol{\Sigma}_k$. Let $\boldsymbol{\mu} = (\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_K)$, and let $\boldsymbol{\Sigma} = (\boldsymbol{\Sigma}_1, \dots, \boldsymbol{\Sigma}_K)$. We can write the model as

$$\mathbf{y}_i | z_i, \boldsymbol{\mu}, \boldsymbol{\Sigma} \sim N(\boldsymbol{\mu}_{z_i}, \boldsymbol{\Sigma}_{z_i}) \quad (1)$$

$$z_i | \boldsymbol{\pi} \sim \text{Multinomial}(\pi_1, \dots, \pi_K). \quad (2)$$

Because of their ability to capture complex distributions, mixture models have been used in applications across a variety of fields (e.g., West and Cao, 1993; Escobar and West, 1995; Müller and Quintana, 2004; Dunson, 2010; Shen et al., 2013).

We use a Bayesian approach to estimate $(\boldsymbol{\mu}, \boldsymbol{\Sigma}, \boldsymbol{\pi})$. As the prior distribution on $\boldsymbol{\pi}$, we use the stick-breaking representation of a truncated Dirichlet process (Ferguson, 1973, 1983; Escobar and West, 1995; West et al., 1994; Sethuraman, 1994; Ishwaran and James, 2001), which has been shown to be convenient for density estimation (Müller and Mitra, 2013). We have

$$\pi_k = v_k \prod_{g < k} (1 - v_g) \quad \text{for } k = 1, \dots, K \quad (3)$$

$$v_k \sim \text{Beta}(1, \alpha) \quad \text{for } k = 1, \dots, K - 1; v_K = 1 \quad (4)$$

$$\alpha \sim \text{Gamma}(a_\alpha, b_\alpha). \quad (5)$$

Following Kim et al. (2014), we set $a_\alpha = b_\alpha = 0.25$ to represent a small effective sample size in the Gamma distribution for α . This allows the likelihood function from the data to dominate the prior distribution for $\boldsymbol{\pi}$. For further discussion about this prior distribution, see Dunson and Xing (2009) and Kim et al. (2014).

For the prior specification of $(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, we consider two alternatives. The most general uses

$$\boldsymbol{\mu}_k | \boldsymbol{\Sigma}_k \sim N(\boldsymbol{\mu}_0, h^{-1} \boldsymbol{\Sigma}_k) \quad (6)$$

$$\boldsymbol{\Sigma}_k \sim \text{InverseWishart}(f, \Phi). \quad (7)$$

Here, f is a prior degrees of freedom and $\Phi = \text{diag}(\phi_1, \dots, \phi_p)$, with $\phi_j \sim \text{Gamma}(a_\phi, b_\phi)$ for $j = 1, \dots, p$. We specify priors with $a_\phi = b_\phi = 0.25$ to allow substantial prior mass at modest sized variances. We set $\boldsymbol{\mu}_0 = \mathbf{0}$, since the variables are standardized, $f = p + 1$ to ensure a proper posterior distribution, and $h = 1$ for convenience. A second alternative, and the one we use in the Census of Manufactures application, replaces (7) with fixed spherical covariance matrices, $\boldsymbol{\Sigma}_k = \sigma \mathbf{I}_p$ for all k and some agency-specified value of $\sigma > 0$. In general, using equal $\boldsymbol{\Sigma}_k$ with a small σ requires more components to describe the distribution adequately than allowing arbitrary $\boldsymbol{\Sigma}_k$. The number of required components increases as σ gets small.

With either specification for $(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, we can use a standard Gibbs sampler to estimate the posterior distribution; see the supplementary material for details. The samplers are easily modified to handle missing values in \mathbf{y}_i that are MAR. Given a current draw of $(\boldsymbol{\mu}, \boldsymbol{\Sigma}, \mathbf{Z})$, we take a draw of the missing values in \mathbf{y}_i from the conditional normal distribution given the observed values in \mathbf{y}_i . We use $N(\boldsymbol{\mu}_{z_i}, \boldsymbol{\Sigma}_{z_i})$ to derive the relevant conditional distribution.

To specify K , we recommend starting with a somewhat large value, for example $K = 30$. In each iteration of the Gibbs sampler, we recommend counting the number of components that include at least one \mathbf{y}_i . If this count ever reaches K , it is prudent to increase K and refit the model with more components. When the count is always less than K , i.e., in every iteration some components are empty, then the value of K is reasonable.

2.2 Pattern mixture modeling

When \mathbf{Y} suffers from unit nonresponse, we can use the model in (1)–(7), possibly with $\Sigma_k = \sigma \mathbf{I}_p$ for all k , to impute values for the $n_0 = n - n_t$ cases with $r_{it} = 0$. At every iteration in the Gibbs sampler, we draw a value of z_i from (2) for each case with $r_{it} = 0$, and draw a value of \mathbf{y}_i from the corresponding $N(\boldsymbol{\mu}_{z_i}, \Sigma_{z_i})$. This generates imputed values from the same distribution as \mathbf{Y}_t ; that is, the unit nonresponse is assumed to be MAR.

To construct pattern mixture models for NMAR unit nonresponse, we again estimate the mixture normal model with \mathbf{Y}_t . In each iteration of the Gibbs sampler, we evaluate the posterior density from (1)–(7). Let $(\boldsymbol{\mu}^m, \Sigma^m)$ equal the values of $(\boldsymbol{\mu}, \Sigma)$ at the iteration with the maximum posterior value (MAP) of the posterior density (Fraley and Raftery, 2007); note that $\Sigma^m = \sigma \mathbf{I}_p$ when using the equal covariance model. Let $\boldsymbol{\pi}^m$ equal the value of $\boldsymbol{\pi}$ at the MAP iteration. Rather than draw z_i for the n_0 nonrespondents using $\boldsymbol{\pi}^m$, which would imply a MAR imputation process, we define and use new mixture probabilities $\boldsymbol{\pi}^* = (\pi_1^*, \dots, \pi_K^*)$. These are chosen to reflect beliefs about the nonrespondents' distribution, as we illustrate in Section 3. We then generate the unit nonrespondents' values from

$$z_i | \boldsymbol{\pi}^*, r_{it} = 0 \sim \text{Multinomial}(\pi_1^*, \dots, \pi_K^*) \quad (8)$$

$$\mathbf{y}_i | z_i, \boldsymbol{\mu}^m, \Sigma^m, r_{it} = 0 \sim N(\boldsymbol{\mu}_{z_i}^m, \Sigma_{z_i}^m). \quad (9)$$

By fixing $(\boldsymbol{\mu}^m, \Sigma^m)$ and altering $\boldsymbol{\pi}^m$, we assume the nonrespondents belong to the same components as \mathbf{Y}_t but at different frequencies. This allows the agency to upweight or downweight regions of support evident in \mathbf{Y}_t when specifying the nonrespondents' distribution. For example, analysts can set π_k^* to be small (or zero) for components with most of the density over regions of \mathbf{Y} that are deemed improbable for nonrespondents. Or, analysts can increase π_k^* for components with density in regions that are deemed more probable than observed in \mathbf{Y}_t . The closer $\boldsymbol{\pi}^*$ is to $\boldsymbol{\pi}^m$, the closer the data are to MAR.

Specifying the entire vector of probabilities can be complicated in some situations, particularly with large p . To facilitate the process of setting $\boldsymbol{\pi}^*$, we developed the NIMC (Nonignorable missingness Imputation for Multivariate Continuous data) tool. The NIMC is an R (R Core Team, 2013) application developed with the `shiny` package (RStudio and Inc., 2014) that allows users to alter components' probabilities via a graphical interface. The NIMC starts at the MAR model, i.e., using $\boldsymbol{\pi}^m$, allowing the user to create and examine departures from MAR by increasing or decreasing any π_k^m . The NIMC renorms the altered probabilities of all components to sum to one, to ensure the resulting mixture model is a proper density. Using the renormed probabilities, the application automatically generates imputations for the missing cases following the model described in (8) and (9), so that the analyst can

change $\boldsymbol{\pi}^*$ to arrive at desired distributions. An illustrative version of the NIMC tool is available at (<http://sites.duke.edu/tcrn/research-projects/downloadable-software/>); we are developing a package to be submitted to CRAN.

By controlling $\boldsymbol{\pi}^*$, the analyst controls where the nonrespondents are likely to fall in the space defined by the respondents. We recommend adjusting ten to fifteen values of π_k^m , and setting the remainder either to zero or to a renormalized version of π_k^m . In our experience, this provides enough flexibility to generate different types of nonrespondent distributions without becoming so large as to make specification of $\boldsymbol{\pi}^*$ too cumbersome.

2.3 Implementing this pattern mixture approach

Before illustrating the approach on the Census of Manufactures data, we discuss some general implementation issues, in particular the choice of unequal or equal $\boldsymbol{\Sigma}_k$, what to do with unoccupied components, steps to facilitating interpretations, and what to do when nonrespondents are unlike any respondents.

For some \mathbf{Y}_t , the model with unequal $\boldsymbol{\Sigma}_k$ results in only a small number of occupied components or results in many components with substantially overlapping regions of high density. For such \mathbf{Y}_t , it can be difficult to alter $\boldsymbol{\pi}^m$ in ways that enable meaningful sensitivity analysis, as it is difficult to generate different distributions for nonrespondents and respondents by adjusting $\boldsymbol{\pi}^m$. In these cases, we recommend setting $\boldsymbol{\Sigma}_k = \sigma \mathbf{I}_p$ for all k , and trying multiple values of σ until finding one that results in fitted components amenable to creating NMAR missing data scenarios. In our experience, when modeling standardized \mathbf{y}_i we find that using $0 < \sigma < 1$ results in usable numbers of occupied components.

When the specified K is large, we expect some components in the MAP iteration to be empty. We suggest setting $\pi_k^* = 0$ for all empty components. The $(\boldsymbol{\mu}_k^m, \boldsymbol{\Sigma}_k^m)$ for empty components is derived entirely from the prior distribution. Such $(\boldsymbol{\mu}_k^m, \boldsymbol{\Sigma}_k^m)$ are not useful for generating nonrespondents' values in a pattern mixture model approach.

To facilitate setting $\boldsymbol{\pi}^*$, it can be useful to rank the components after finding $(\boldsymbol{\mu}^m, \boldsymbol{\Sigma}^m, \boldsymbol{\pi}^m)$. In the Census of Manufactures application, we rank components based on their distance to an artificially created point with small values of all variables. Specifically, we rank on $\Delta_k = (\boldsymbol{\mu}_k - \mathbf{y}_{\min})'(\boldsymbol{\mu}_k - \mathbf{y}_{\min})$, where $\mathbf{y}_{\min} = (\min_i(y_{11}, \dots, y_{n1}), \dots, \min_i(y_{1p}, \dots, y_{np}))$. Rankings facilitate interpretation of the components. For example, to impute more nonrespondents' data with large values of \mathbf{y}_i , we inflate the π_k^m of the components with the highest ranks on Δ_k . To generate nonrespondents' data with low values of \mathbf{y}_i we inflate the π_k^m of the components with lowest ranks on Δ_k .

Even though changing $\boldsymbol{\pi}^m$ results in different respondents' and nonrespondents' distributions, the nonrespondents' values are imputed in the same regions as \mathbf{Y}_t , albeit with different propensities. Adjusting $\boldsymbol{\pi}^m$ does not capture NMAR settings where the missing values are suspected to be in regions with no respondents' data. To represent such scenarios, the analyst must create new components. This requires specifying new probabilities and the mean and covariance matrix of the new components.

3 Imputation of Nonignorable Unit Nonresponse in the Census of Manufactures

We now illustrate the proposed approach to pattern mixture modeling using data from the 2007 U. S. Census of Manufactures (CMF). The CMF covers all establishments in manufacturing industries in the United States. Data are collected via mail-in questionnaires (U. S. Census Bureau, 2015). We focus on three highly important variables in the CMF, namely the total value of shipments (TVS), total employment (TE), and salary/wages (SW). We use data from the ready-mix concrete industry, which is relatively homogeneous in these three variables as most of their business is based on a single product. We also examined the plastics products manufacturing industry, which is a more broad classification and more heterogeneous. The results for the plastics products industry are included in the supplementary material.

CMF data are collected and processed in waves during the year, depending on when each establishment mails its form. This can result in NMAR unit nonresponse at some waves. For example, larger establishments may employ individuals who are intimately familiar with and responsible for the types of variables requested in the CMF, so that it is straightforward for them to return the questionnaire. On the other hand, smaller companies may not have such employees, so that it takes them longer to get the right data and return the completed CMF forms. The CMF data include response date, so that we can tell when establishments returned their questionnaires, if at all.

For illustration, we consider as \mathbf{Y}_t all the establishments that had at least one reported value among the three variables of interest and that had a valid response date. We consider the remaining establishments as unit nonrespondents. We note that some establishments in the file have data on the three key variables but no response date. We consider these cases as unit nonrespondents and impute their entire vector of survey variables. Due to disclosure limitations imposed by the Census Bureau, we cannot provide the sample sizes.

We consider three scenarios when setting $\boldsymbol{\pi}^*$. The first is to assume MAR nonresponse, in which we set $\boldsymbol{\pi}^* = \boldsymbol{\pi}^m$. The second scenario assumes that nonrespondents tend to be small establishments. To generate nonrespondents' data consistent with this scenario, we inflate π_k^m for components with small values of Δ_k . The third scenario represents an opposite NMAR mechanism: we inflate π_k^m for components with large values of Δ_k . This scenario is in some ways a worst case scenario, since large establishments can have disproportionate impacts on summary statistics (e.g., totals) typically computed with CMF data. For all scenarios, we use the mixture normal model with $\boldsymbol{\Sigma}_k = \sigma_p \mathbf{I}_p$ for all k , where $\sigma = 0.3$. Based on investigations with different values of σ , the value 0.3 offers sufficient flexibility to impute data in all regions of the variable space while being large enough to make choosing $\boldsymbol{\pi}^*$ not cumbersome. The MAP iteration has 17 non-empty components.

Because the data are heavily skewed, we log transform and standardize the variables. We exclude the few cases with zero value for shipments, employment or salary when estimating the model (alternatively, one could replace the zeros with small positive values). Figure 1 displays pairwise scatter plots for the concrete industry after imputation using MAR. Most

of the density is concentrated in fewer than half of the components. As desired in the MAR scenario, the imputed values for nonresponse have essentially the same distribution as \mathbf{Y}_t , which suggests that the mixture normal model describes the distribution of \mathbf{Y}_t reasonably well.

Figure 2 and Figure 3 display results for the concrete industry after imputation for the two NMAR scenarios. In Figure 2, the imputed values for nonrespondents are primarily in regions with small values of \mathbf{Y}_t . We accomplish this by setting π_k^m to be large for $k \in (11, \dots, 16)$ (the low-ranking components), taking mass away from components in the middle and upper tails of the distribution, e.g., $k \in (4, 6)$. We note that the NMAR mechanism still generates large values in the imputed \mathbf{y}_i , just not as frequently as would be expected under MAR. This is because we allow $\pi_k^* > 0$ for all occupied components, as well as the random chance of imputing values in regions of low density within low-ranking components.

Similarly, in Figure 3 we set π_k^* to be large for the top-ranking components, generating imputed values for nonrespondents that tend to be in the region of high values of \mathbf{Y}_t . We defer analysis of the impact of these assumptions on inferences until we present the approach for evaluating stopping decisions.

4 Using the NMAR Imputations in Stopping Rules

At any time t , the agency needs to decide how many, if any, of the $n_0 = n - n_t$ nonrespondents to target for follow up. In this section, we describe how the NMAR imputation approach can be used to inform such decisions. The basic idea is to create completed datasets under various NMAR scenarios, take proposed nonresponse followup samples from those completed datasets, and evaluate the cost and accuracy associated with different sample sizes.

4.1 Creating completed datasets

Let \mathbf{Y}_o be the data for the n_0 unit nonrespondents at time t . As these data are unknown, we begin by creating completed versions of $\mathbf{Y} = (\mathbf{Y}_t, \mathbf{Y}_o)$. These serve as plausible “true values” that facilitate evaluation of agency nonresponse followup decisions. We follow the approach of Section 2.2 to create the completed versions, in which we (i) fit the multivariate mixture model using \mathbf{Y}_t and (ii) generate several, say S , imputation scenarios by adjusting $\boldsymbol{\pi}^*$. Let $\boldsymbol{\pi}^{*s}$ be the adjusted $\boldsymbol{\pi}$ for missingness scenario s , where $s = 1, \dots, S$. For each $\boldsymbol{\pi}^{*s}$, we generate M hypothetical true datasets, $\mathbf{Y}^s = (\mathbf{Y}^{s,1}, \dots, \mathbf{Y}^{s,M})$, where $\mathbf{Y}^{s,j} = (\mathbf{Y}_t, \mathbf{Y}_0^{s,j})$ and $\mathbf{Y}_0^{s,j}$ for $j = 1, \dots, M$ is an independent draw from (8) and (9) with $\boldsymbol{\pi}^* = \boldsymbol{\pi}^{*s}$. The use of $S > 1$ scenarios allows the agency to assess the implications of nonresponse followup decisions under different plausible mechanisms for unit nonresponse. The use of $M > 1$ completed datasets allows the agency to propagate uncertainty due to the unit nonresponse and imputation procedure. For simplicity, we impute any missing values in \mathbf{Y}_t due to item nonresponse using $\boldsymbol{\pi}^m$ and the parameters from the MAP iteration.

Let $n_{max} \leq n_0$ be the maximum number of nonrespondents that hypothetically could be collected given the total budget and time available for nonresponse followup. Let $n_f = \delta n_{max}$,

where $\delta \in [0, 1]$, be a number of nonrespondents that the agency is considering to include in a followup sample at time t . Here, $n_f = 0$ means that the agency stops data collection at time t . The agency takes an independent random sample of n_f (possibly zero) individuals from each $\mathbf{Y}^{s,j}$. These represent hypothetical nonresponse followup samples. After followup, the data observed by the agency would include $(\mathbf{Y}_t, \mathbf{Y}_f^{s,j})$, where $\mathbf{Y}_f^{s,j}$ is the sampled nonrespondents' values. For now, we assume no additional nonresponse in the followup sample; we discuss relaxations of this Section 6.

Let $\mathbf{Y}_{0f}^{s,j}$ be the data for the $n_{0f} = n_0 - (n_t + n_f)$ individuals that still would be unobserved by the agency after the followup sampling. As $\mathbf{Y}_{0f}^{s,j}$ would be unknown, following the multiple imputation framework we need to impute plausible values for these n_{0f} cases, thereby creating completed versions of \mathbf{Y} . To do so, we assume the remaining nonresponse is MAR (even if in reality it is not), which we believe mimics what most agencies would do when faced with imputing nonrespondents' values after completing all followup activities. We fit the mixture model again using only the observations in $\mathbf{Y}_0^{s,j}$, as depicted in Figure 4(a), and create R completed datasets by drawing new versions of $\mathbf{Y}_{0f}^{s,j}$ from the corresponding posterior predictive distribution. Estimating the imputation model with only $\mathbf{Y}_0^{s,j}$ corresponds to the belief that the nonsampled nonrespondents are more similar to the individuals in the latest wave of nonresponse followup than to the individuals in \mathbf{Y}_t . When n_f is too small to support reliable modeling, agencies may choose to fit the mixture model to $(\mathbf{Y}_t, \mathbf{Y}_f^{s,j})$, as shown in Figure 4(b).

In sum, for each value of $\boldsymbol{\pi}^{*s}$ where $s = 1, \dots, S$, we use the following procedure to generate completed versions of \mathbf{Y} .

1. Generate M datasets $(\mathbf{Y}^{(s,1)}, \dots, \mathbf{Y}^{(s,M)})$ using independent draws of the nonrespondents' values based on $\boldsymbol{\pi}^{*s}$ and the mixture model estimated on \mathbf{Y}_t . These represent hypothetical "true values" of \mathbf{Y} .
2. For the value of δ under consideration and $j = 1, \dots, M$, repeat the following steps.
 - (a) Take a random sample of n_f individuals from each $\mathbf{Y}_0^{(s,j)}$, i.e., from the nonrespondents' "true values." Let $\mathbf{Y}_f^{(s,j)}$ denote this follow-up sample set.
 - (b) Fit a new mixture model to $\mathbf{Y}_f^{(s,j)}$ provided n_f is large enough to support reliable modeling. If not, fit a new mixture model to $(\mathbf{Y}_t, \mathbf{Y}_f^{(s,j)})$.
 - (c) Generate R multiply imputed values for the n_{0f} observations that remain nonrespondents, using the posterior predictive distribution of the new mixture model and assuming MAR. For $l = 1, \dots, R$, let $\mathbf{Y}_{0f}^{(s,j,l)}$ be the l th draw of the remaining nonrespondents' values, and let $\mathbf{Y}_\delta^{(s,j,l)} = (\mathbf{Y}_t, \mathbf{Y}_f^{(s,j)}, \mathbf{Y}_{0f}^{(s,j,l)})$ be a completed version of \mathbf{Y} for the specified δ .

The agency can repeat Step 2 for multiple values of δ , each time using the same $\{\mathbf{Y}^{(s,j)} : s = 1, \dots, S; j = 1, \dots, m\}$. The agency then runs key analyses on $\{\mathbf{Y}^{(s,j)} : s = 1, \dots, S; j = 1, \dots, M\}$ and on $\{\mathbf{Y}^{(s,j,l)} : s = 1, \dots, S; j = 1, \dots, M; l = 1, \dots, R\}$ and evaluates the differences among them, as we describe next.

4.2 Utility measures

The decision of stopping the data collection depends critically on the amount of information that has been collected so far, and the amount of information that has yet to be collected. For example, when the unit nonresponse at some time t is close to MAR and n_t is already large, collecting more data may not improve accuracy substantially. On the other hand, when the unit nonresponse is NMAR and $n - n_t$ is large, not sampling the nonrespondents could sacrifice accuracy too much. We propose three measures for assessing accuracy of followup sampling choices for multivariate continuous data. The measures do not consider survey weights; we discuss extensions for complex designs in Section 6.

4.2.1 Relative percent differences

For $s = 1, \dots, S$, $j = 1, \dots, M$, and $v = 1, \dots, p$, let $\bar{y}_v^{(s,j)} = \sum_{i=1}^n y_{i,v}^{(s,j)} / n$ be the marginal mean of the variable indexed by v computed with $\mathbf{Y}^{(s,j)}$. Similarly, let $\bar{y}_{v,\delta}^{(s,j,l)} = \sum_{i=1}^n y_{i,v}^{(s,j,l)} / n$ denote the marginal mean of the variable indexed by v computed with $\mathbf{Y}_\delta^{(s,j,l)}$. For any s and v , we compute the average percentage difference,

$$\theta_{v,\delta}^{(s)} = \frac{1}{MR} \sum_{j=1}^M \sum_{l=1}^R \left| \frac{\bar{y}_v^{(s,j)} - \bar{y}_{v,\delta}^{(s,j,l)}}{\bar{y}_v^{(s,j)}} \right|. \quad (10)$$

Large values of $\theta_{v,\delta}^{(s)}$ are less desirable than small values. When values of $\theta_{v,\delta}^{(s)}$ do not change markedly as n_f increases, higher intensity nonresponse followup sampling does not return significant gains in accuracy for variable v . We note that one could replace means with any quantity of substantive relevance. As a one number summary, we compute $\theta_\delta^{(s)} = \sum_{v=1}^p \theta_{v,\delta}^{(s)} / p$.

4.2.2 Difference as fraction of standard error

For $s = 1, \dots, S$, $j = 1, \dots, M$, and $v = 1, \dots, p$, let $\hat{\sigma}_v^{(s,j)}$ and $\hat{\sigma}_{v,\delta}^{(s,j,l)}$ denote the empirical standard deviations for variable v computed with $\mathbf{Y}^{(s,j)}$ and $\mathbf{Y}_\delta^{(s,j,l)}$, respectively. To compute each $\hat{\sigma}_v$ we use the square root of the usual formula, $\sum_{i=1}^n (y_{i,v} - \bar{y}_v)^2 / (n-1)$. For any (s, j, l) and v , we compute the average difference in the means as a fraction of an approximate standard error,

$$\tau_{v,\delta}^{(s,j,l)} = \frac{\bar{y}_v^{(s,j)} - \bar{y}_{v,\delta}^{(s,j,l)}}{\sqrt{\left[\left(\hat{\sigma}_v^{(s,j)} \right)^2 + \left(\hat{\sigma}_{v,\delta}^{(s,j,l)} \right)^2 \right] / 2}}. \quad (11)$$

For any s and v , we compute

$$\tau_{v,\delta}^{(s)} = \frac{1}{MR} \sum_{j=1}^M \sum_{l=1}^R \left| \tau_{v,\delta}^{(s,j,l)} \right|, \quad (12)$$

Like $\theta_{v,\delta}^{(s)}$, this criterion penalizes large differences in the means of $\mathbf{Y}^{s,j}$ and the completed datasets. However, it considers differences in terms of standardized units. As a one number summary, we compute $\tau_\delta^{(s)} = \sum_{v=1}^p \tau_{v,\delta}^{(s)} / p$.

4.2.3 Propensity score distances

Woo et al. (2009) present a metric for quantifying the similarity of the multivariate empirical distributions in two datasets. They develop this metric to evaluate the analytic validity of redacted data, comparing the empirical distributions in the redacted and corresponding confidential data. We adapt their metric for nonresponse followup decisions as follows. For each (s, j, l) , we concatenate $\mathbf{Y}^{(s,j)}$ and $\mathbf{Y}_\delta^{(s,j,l)}$. We add a variable $T_i = 1$ for all rows belonging to $\mathbf{Y}^{(s,j)}$ and $T_i = 0$ for all rows belonging to $\mathbf{Y}_\delta^{(s,j,l)}$. We then regress T on main effects for (y_1, \dots, y_p) using a generalized additive model estimated on the concatenated data (Woo et al., 2008). With this regression, we compute the predicted probability for each observation, $\hat{e}_{i,\delta}^{(s,j,l)}$. For any (s, j, l) , we compute

$$\rho_\delta^{(s,j,l)} = \frac{\sum_{i=1}^{2n} (\hat{e}_{i,\delta}^{(s,j,l)} - 0.5)^2}{2n}. \quad (13)$$

When $\mathbf{Y}^{(s,j)}$ and $\mathbf{Y}_\delta^{(s,j,l)}$ have similar empirical distributions, the values of $\hat{e}_{i,\delta}^{(s,j,l)}$ should be close to 0.5, so that $\rho_\delta^{(s,j,l)}$ should be near zero (Woo et al., 2009).

As a metric, we compute the average,

$$\rho_\delta^{(s)} = \frac{1}{MR} \sum_{j=1}^M \sum_{l=1}^R \rho_\delta^{(s,j,l)}. \quad (14)$$

When $\rho_\delta^{(s)}$ is close to zero, the “true value” and multiply-imputed databases contain similar information content; in this case, nonresponse followup may not be worth the cost. When ρ is close to 1/4, the “true value” and multiply-imputed databases have quite different distributions, suggesting that there is information to be gained by additional followup sampling. Woo et al. (2009) do not suggest values of ρ that can be considered “good enough.” Thus, we recommend comparing relative values of ρ for different follow-up sample sizes and scenarios.

4.3 Cost measure

As n_f increases, $\mathbf{Y}^{(s,j)}$ and corresponding $\mathbf{Y}_\delta^{(s,j,l)}$ have increasing numbers of overlapping records, causing the measures of Section 4.2 to decrease. Of course, agencies must balance accuracy gains with increased expenses. Therefore, we inject cost considerations into the decision process, so that agencies can decide how much they are willing to spend for utility improvements.

For any (s, j) , let $C_f^{(s,j)}$ be the total cost of the followup sample, $C_0^{(s,j)}$ be the fixed cost regardless of the sample size, and $c_i^{(s,j)}$ the cost of selecting, measuring and processing unit

i. We assume the cost function,

$$C_f^{(s,j)} = C_0^{(s,j)} + \sum_{i \in \mathbf{Y}_f^{s,j}} c_i^{(s,j)}. \quad (15)$$

For the illustration with CMF data, we assume that $c_i^{(s,j)} = c$ for all nonrespondents and all (s, j) ; that is, we make nonresponse followup costs a linear function of n_f . We also assume that $C_0^{(s,j)} = C_0$ for all (s, j) . Under these assumptions, n_{max} is the largest value of n_f such that $C_0 + cn_f$ is less than the total budget available for nonresponse followup at time t .

The framework readily handles more complex cost functions (for examples, see Groves, 2004). When agencies allow $c_i^{(s,j)}$ to vary across individuals or scenarios, one can add in (15) the realized cost for each individual sampled in Step 2(a) of the algorithm from Section 4.

5 Nonresponse Followup Decisions in the Census of Manufactures Data

The U. S. Census Bureau spends significant resources on nonresponse followup operations during the CMF data collection. For example, they mail additional questionnaires to nonrespondents and designate Census consultants to work with some of the largest establishments. The entire data collection process takes more than a year to complete. Thus, the Census Bureau potentially could benefit from sensible stopping decisions that reduce costs and shorten time to data release without too much sacrifice in accuracy.

To illustrate the methodology, we use the data and missingness scenarios described in Section 3. For brevity, we report only results for the concrete industry. The results for the plastics products industry can be found in the supplementary material. We consider $\delta \in (0, 0.25, 0.5, 0.75, 1)$, that is, collecting follow-up samples with 0%, 25%, 50%, 75% and 100% of n_{max} . We set $n_{max} = n_0$, corresponding to sufficient resources to attempt to collect all nonrespondents' data if desired. For each scenario s , we create $M = 10$ completed hypothetical censuses in the two industries. For each $\mathbf{Y}^{(s,j)}$ and for each value of δ , we take simple random followup samples to obtain hypothetical versions of $\mathbf{Y}_{f,\delta}^{(s,j)}$. We generate $R = 5$ completed datasets $\mathbf{Y}_\delta^{(s,j,l)}$ by fitting new mixture models to $\mathbf{Y}_{f,\delta}^{(s,j)}$ only; for comparison, we also report results with mixture models based on $(\mathbf{Y}_t, \mathbf{Y}_{f,\delta}^{(s,j)})$. Due to disclosure limitation requirements, at the request of the Census Bureau we do not report values for the cost measure. Since the cost measure is a linear function of n_f , we believe that we can illustrate the decision-making process and sensitivity analysis using results based on δ .

Figure 5 displays $(\theta_\delta^{(s)}, \tau_\delta^{(s)}, \rho_\delta^{(s)})$ for the three missingness scenarios in the concrete industry data as a function of δ . As expected, in all cases the measures decrease as δ increases. When $\delta = 1$, all units are sampled so that all utility measures equal zero. When $\delta = 0$, imputations are only based on \mathbf{Y}_t . In the MAR scenario, we produce greater accuracy when generating the $R = 5$ imputed datasets from the model estimated with $(\mathbf{Y}_t, \mathbf{Y}_{f,\delta}^{(s,j)})$ than

the model estimated with \mathbf{Y}_t . This is expected, since the former model is based on more observations than the latter, and both \mathbf{Y}_t and $\mathbf{Y}_{f,\delta}^{(s,j)}$ are representative of the nonrespondents under MAR. In contrast and also as expected, in the NMAR scenarios the values of the utility measures are much smaller when the imputation models are estimated only with $\mathbf{Y}_{f,\delta}^{(s,j)}$, which is representative of the nonrespondents whereas \mathbf{Y}_t is not. The values of the utility measures are generally smaller for the MAR scenario than for the NMAR scenario. The values in the MAR scenario can serve as baselines, providing a sense of the magnitudes of “best case” results.

Figure 5 can be used to make qualitative decisions about nonresponse followup. When the agency suspects the nonrespondents’ data are close to MAR, conducting a followup sample does not improve accuracy dramatically with respect to these measures. For example, a 25% nonresponse followup sample decreases $\theta_\delta^{(s)}$ by around 7 points (from .255 to .185), and a 50% nonresponse followup sample decreases $\theta_\delta^{(s)}$ by around 11.6 points (from .255 to .139). In contrast, when the agency is concerned that the data are NMAR in ways defined by the two selected $\boldsymbol{\pi}^{*s}$, collecting data from 25% or 50% of the nonrespondents results in massive accuracy improvements: $\theta_\delta^{(s)}$ drops from values of at least 1.0 when $n_f = 0$ to values in line with those for the MAR scenario. For all scenarios, beyond $\delta = .50$ the gains seem unlikely to be worth the cost. We reach similar qualitative conclusions for all three utility measures.

Agencies may want to plug utility and cost measures into specific loss functions, and find n_f that minimize loss. We leave specification of suitable loss functions to future work.




Figure 1: Pairwise scatter plots for the concrete industry from the 2007 CMF for the MAR imputation scenario. Observed points are plotted as gray hollow circles. The black filled circles on the lower diagonal are the imputed points. The colored circles on the upper diagonal are the 95% quantile ellipses of the fitted clusters, with color intensity proportional to the mixture probabilities. Graph copied from the output of the NIMC R tool. Axes removed to prevent disclosure of information about the magnitude of the data.




Figure 2: Pairwise scatter plots for the concrete industry from the 2007 CMF for the NMAR imputation scenario with higher probabilities for bottom ranked clusters. Observed points are plotted as gray hollow circles. The black filled circles on the lower diagonal are the imputed points. The colored circles on the upper diagonal are the 95% quantile ellipses of the fitted clusters, with color intensity proportional to the mixture probabilities. Graph copied from the output of the NIMC R tool. Axes removed to prevent disclosure of information about the magnitude of the data.




Figure 3: Pairwise scatter plots for the concrete industry from the 2007 CMF for the NMAR imputation scenario with higher probabilities for top ranked clusters. Observed points are plotted as gray hollow circles. The black filled circles on the lower diagonal are the imputed points. The colored circles on the upper diagonal are the 95% quantile ellipses of the fitted clusters, with color intensity proportional to the mixture probabilities. Graph copied from the output of the NIMC R tool. Axes removed to prevent disclosure of information about the magnitude of the data.




Figure 4: Two possible approaches to estimating the imputation model after hypothetical nonresponse followup sampling. The left panel uses only the latest wave of nonrespondents, and the right panel uses all collected data up to time t .




Figure 5: Summary of utility measures for the three scenarios for the concrete industry from the 2007 CMF. The results for ρ , τ and θ are on the first, second and last row, respectively. The plots on the right column display the results with the model fit to observed data and follow-up sample, and the plots on the right display the results with the model fit to follow-up sample only. The faded points are the individual values for each multiple imputation.

6 Concluding Remarks

It is prudent for agencies to consider several plausible mechanisms for unit nonresponse when assessing sensitivity of inferences to NMAR nonresponse and making decisions about followup. Agencies can ease interpretations by using expected values of the utility and cost measures. To compute these, agencies can put a subjective probability p_s on each missing data scenario s under consideration, representing the agency's belief on the chance of each scenario being (close to) reality. The agency can average the utility measures and costs over the subjective probabilities, e.g., $\sum_s \theta_\delta^{(s)} p_s$.

In any followup sample, it is likely that some sampled individuals will continue to be nonrespondents. To account for this in the decision-making process, agencies can propose models for the nonresponse in the followup samples. For example, the agency can randomly sample nonrespondents within the followup sample from a specified selection model. To mimic typical practice and simplify the imputation process, the agency could combine any nonrespondents in the followup sample with the nonsampled cases, and use MAR models based on respondents in $\mathbf{Y}_{f,\delta}^{(s,j)}$ to create the completed $\mathbf{Y}_\delta^{(s,j,l)}$. Investigating how such additional modeling affects decisions about n_f is a topic for future research.

We present the approach for deciding nonresponse followup sample size for a single time point. However, agencies could adapt the methodology for multiple decision points. For example, in the CMF the Census Bureau collects and processes data in waves during the year, so that it has the opportunity to adapt nonresponse followup designs at multiple time points, e.g., monthly or quarterly. The agency can reset the process at each t , using the collected data up to that time as \mathbf{Y}_t and recreating hypothetical nonrespondents' data.

In some databases the agency may possess external information on nonrespondents, such as paradata or administrative records. In the CMF, data on establishments' employment, sales, and payroll are available to the Census Bureau from the U. S. Internal Revenue Service. When these administrative values are highly correlated with the recorded values, the agency could use the auxiliary data to improve imputation, for example by including them in the mixture model as additional variables or as regressors in the component means. In such cases, it may be reasonable to impute missing values (for establishments and variables with administrative data) under MAR assumptions. Agencies still may want to assess sensitivity to nonignorable missingness mechanisms when using auxiliary variables. This can be done using the approach described in Section 2.

The methodology is directly applicable to census data, but it could be modified for surveys. The key issue is to model relationships between the survey design variables and the substantive survey variables. We conjecture that one could regress on functions of the design variables when specifying the means of the mixture components. Alternatively, a more *ad hoc* means to do so is to include the survey weights as a variable in the mixture model.

7 Acknowledgments

Any opinions and conclusions expressed in this article are those of the authors and do not necessarily represent the views of the U. S. Census Bureau. All results have been reviewed to ensure that no confidential information is disclosed.

References

- Daniels, M. and Hogan, J. (2008). *Missing Data in Longitudinal Studies: Strategies for Bayesian Modeling and Sensitivity Analysis*. Chapman & Hall/CRC Press. Taylor & Francis, Boca Raton.
- Diggle, P. and Kenward, M. G. (1994). Informative drop-out in longitudinal data analysis. *Applied statistics*, 43(1):49–93.
- Dunson, D. B. (2010). Nonparametric Bayes applications to biostatistics. In Hjort, N., Holmes, C., Müller, P., and Walker, S., editors, *Bayesian Nonparametrics*, Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press.
- Dunson, D. B. and Xing, C. (2009). Nonparametric Bayes modeling of multivariate categorical data. *Journal of the American Statistical Association*, 104(487):1042–1051.
- Escobar, M. D. and West, M. (1995). Bayesian density estimation and inference using mixtures. *Journal of the American Statistical Association*, 90(430):577–588.
- Ferguson, T. S. (1973). A Bayesian analysis of some nonparametric problems. *The Annals of Statistics*, 1(2):209–230.
- Ferguson, T. S. (1983). Bayesian density estimation by mixtures of normal distributions. In *Recent advances in statistics*, volume 24, pages 287–302. New York Academic Press.
- Finamore, J., Reist, B., and Coffey, S. (2013). 2013 National Survey of College Graduates: A practice-based investigation of adaptive design. In *AAPOR (American Association for Public Opinion Research) 68th Annual Conference*.
- Fraley, C. and Raftery, A. E. (2007). Bayesian regularization for normal mixture estimation and model-based clustering. *Journal of Classification*, 24(2):155–181.
- Ghosal, S. and van der Vaart, A. (2001). Entropies and rates of convergence for maximum likelihood and Bayes estimation for mixtures of normal densities. *The Annals of Statistics*, 29(5):1233–1263.
- Ghosal, S. and van der Vaart, A. (2007). Posterior convergence rates of Dirichlet mixtures at smooth densities. *The Annals of Statistics*, 35(2):697–723.

- Glynn, R., Laird, N., and Rubin, D. (1986). Selection modeling versus mixture modeling with nonignorable nonresponse. In Wainer, H., editor, *Drawing Inferences from Self-Selected Samples*, pages 115–142. Springer New York.
- Glynn, R. J., Laird, N. M., and Rubin, D. B. (1993). Multiple imputation in mixture models for nonignorable nonresponse with follow-ups. *Journal of the American Statistical Association*, 88(423):984–993.
- Groves, R. (2004). *Survey Errors and Survey Costs*. Wiley Series in Probability and Statistics. Wiley, New York.
- Ishwaran, H. and James, L. F. (2001). Gibbs sampling methods for stick-breaking priors. *Journal of the American Statistical Association*, 96(453).
- Kim, H. J., Reiter, J. P., Wang, Q., Cox, L. H., and Karr, A. F. (2014). Multiple imputation of missing or faulty values under linear constraints. *Journal of Business and Economic Statistics*, 32(3):375–386.
- Little, R. (1993). Pattern-mixture models for multivariate incomplete data. *Journal of the American Statistical Association*, 88(421):125–134.
- Little, R. (1994). A class of pattern-mixture models for normal incomplete data. *Biometrika*, 81(3):471–483.
- Little, R. (1995). Modeling the drop-out mechanism in repeated-measures studies. *Journal of the American Statistical Association*, 90(431):1112–1121.
- Little, R. (2008). Selection and pattern-mixture models. In Fitzmaurice, G., Davidian, M., Verbeke, G., and Molenberghs, G., editors, *Longitudinal Data Analysis*, pages 409–431. Chapman and Hall/CRC, Boca Raton.
- Little, R. and Rubin, D. (2002). *Statistical Analysis with Missing Data*. Wiley Series in Probability and Statistics. Wiley, Hoboken, NJ, Second edition.
- Miller, P. V. (2013). Adaptive design at the Census Bureau - a new way of doing business. In *AAPOR (American Association for Public Opinion Research) 68th Annual Conference*, page 171.
- Molenberghs, G., Kenward, M. G., and Lesaffre, E. (1997). The analysis of longitudinal ordinal data with nonrandom drop-out. *Biometrika*, 84(1):33–44.
- Molenberghs, G., Michiels, B., Kenward, M. G., and Diggle, P. J. (1998). Monotone missing data and pattern-mixture models. *Statistica Neerlandica*, 52(2):153–161.
- Müller, P. and Mitra, R. (2013). Bayesian nonparametric inference – why and how. *Bayesian Analysis*, 8(2):269–302.

- Müller, P. and Quintana, F. A. (2004). Nonparametric Bayesian data analysis. *Statistical Science*, 19(1):95–110.
- R Core Team (2013). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.
- Rao, R. S., Glickman, M. E., and Glynn, R. J. (2008). Stopping rules for surveys with multiple waves of nonrespondent follow-up. *Statistics in Medicine*, 27(12):2196–2213.
- RStudio and Inc. (2014). *shiny: Web Application Framework for R*. R package version 0.9.1.
- Rubin, D. (1987). *Multiple Imputation for Nonresponse in Surveys*. Wiley, Hoboken, NJ.
- Rubin, D. B. (1976). Inference and missing data. *Biometrika*, 63(3):581–592.
- Schouten, B., Bethlehem, J., Beullens, K., Kleven, Ø., Loosveldt, G., Luiten, A., Rutar, K., Shlomo, N., and Skinner, C. (2012). Evaluating, comparing, monitoring, and improving representativeness of survey response through r-indicators and partial r-indicators. *International Statistical Review*, 80(3):382–399.
- Schouten, B., Cobben, F., and Bethlehem, J. (2009). Indicators for the representativeness of survey response. *Survey Methodology*, 35(1):101–113.
- Schouten, B., Shlomo, N., and Skinner, C. (2011). Indicators for monitoring and improving representativeness of response. *Journal of Official Statistics*, 27(2):1–24.
- Sethuraman, J. (1994). A constructive definition of Dirichlet priors. *Statistica Sinica*, 4:639–650.
- Shen, W., Tokdar, S. T., and Ghosal, S. (2013). Adaptive Bayesian multivariate density estimation with Dirichlet mixtures. *Biometrika*, 100(3):623–640.
- Thijs, H., Molenberghs, G., Michiels, B., Verbeke, G., and Curran, D. (2002). Strategies to fit pattern-mixture models. *Biostatistics*, 3(2):245–265.
- U. S. Census Bureau (2015). Annual Survey of Manufactures (ASM). <http://www.census.gov/manufacturing/asm/index.html>.
- Wagner, J. (2008). *Adaptive Survey Design to Reduce Nonresponse Bias*. PhD thesis, University of Michigan.
- Wagner, J. and Raghunathan, T. E. (2010). A new stopping rule for surveys. *Statistics in Medicine*, 29(9):1014–1024.
- West, M. and Cao, G. (1993). Assessing mechanisms of neural synaptic activity. In *Case Studies in Bayesian Statistics, Vol. 1*, pages 416–428. Springer-Verlag, New York.

- West, M., Müller, P., and Escobar, M. D. (1994). Hierarchical priors and mixture models, with application in regression and density estimation. In Smith, A. and Freeman, P., editors, *Aspects of Uncertainty: A Tribute to D.V. Lindley*, pages 363–386. London: Wiley.
- Woo, M.-J., Reiter, J. P., and Karr, A. F. (2008). Estimation of propensity scores using generalized additive models. *Statistics in Medicine*, 27(19):3805–3816.
- Woo, M.-J., Reiter, J. P., Oganian, A., and Karr, A. F. (2009). Global measures of data utility for microdata masked for disclosure limitation. *Journal of Privacy and Confidentiality*, 1(1):7.

Supplementary Material for Stop or Continue Data Collection: A Nonignorable Missing Data Approach for Continuous Variables

Thais Paiva¹ and Jerome P. Reiter²

¹Department of Computer Science, Federal University of Minas Gerais, Belo Horizonte, Brazil

²Department of Statistical Science, Duke University, Durham, NC, USA

November 9, 2015

1 Introduction

This document includes supplementary material for the article “Stop or Continue Data Collection: A Nonignorable Missing Data Approach for Continuous Variables”. Section 2 includes the steps and full conditionals for the Gibbs sampler for the mixture model proposed in the main text. Section 3 presents the results for the plastics products industry from the Census of Manufactures data.

2 Gibbs Sampler for the Mixture of Multivariate Normal Distributions

With the conjugate prior specification from the model (??)–(??) described in Section ?? from the main text, we obtain posterior samples using a Gibbs sampling algorithm (Ishwaran and James, 2001) that we now describe.

For each component $k = 1, \dots, K$, let $N_k = \sum_{i=1}^n \mathbb{1}_{(z_i=k)}$ be the number of observations in component k at a particular iteration of the sampler. Let $\bar{\mathbf{y}}_k = \sum_{\{i:z_i=k\}} \mathbf{y}_i / N_k$ and $S_k = \sum_{\{i:z_i=k\}} (\mathbf{y}_i - \bar{\mathbf{y}}_k)(\mathbf{y}_i - \bar{\mathbf{y}}_k)'$. At each iteration, we recompute each $(N_k, \bar{\mathbf{y}}_k, S_k)$. After initialization, the sampler iterates through the following steps.

1. For $k = 1, \dots, K$, update $\boldsymbol{\mu}_k$ and Σ_k from the full conditionals,

$$\Sigma_k | \mathbf{y}, \mathbf{z} \sim \text{InverseWishart}(f_k, \Phi_k) \quad (1)$$

$$\boldsymbol{\mu}_k | \Sigma_k, \mathbf{y}, \mathbf{z} \sim N(\bar{\boldsymbol{\mu}}_k, \bar{\Sigma}_k), \quad (2)$$

where $f_k = f + N_k$ and $\Phi_k = \Phi + S_k + (\bar{\mathbf{y}}_k - \boldsymbol{\mu}_0)(\bar{\mathbf{y}}_k - \boldsymbol{\mu}_0)'/(1/h + 1/N_k)$ are the parameters for sampling the covariance matrix, $\bar{\boldsymbol{\mu}}_k = (h\boldsymbol{\mu}_0 + N_k\bar{\mathbf{y}}_k)/(h + N_k)$, and $\bar{\Sigma}_k = \frac{1}{h+N_k}\Sigma_k$.

2. For $k = 1, \dots, K - 1$, update v_k from the full conditional,

$$v_k | \mathbf{z}, \alpha \sim \text{Beta} \left(1 + N_k, \alpha + \sum_{g>k} N_g \right), \quad (3)$$

and set $v_K = 1$. For $k = 1, \dots, K$, set $\pi_k = v_k \prod_{g<k} (1 - v_g)$.

3. For $j = 1, \dots, p$, update ϕ_j from the full conditional,

$$\phi_j | \Sigma \sim \text{Gamma} \left(a_\phi + \frac{K(p+1)}{2}, b_\phi + \frac{1}{2} \sum_{k=1}^K (\Sigma_k^{-1})_{(j,j)} \right), \quad (4)$$

where $(\Sigma_k^{-1})_{(j,j)}$ is the j -th diagonal element of Σ_k^{-1} .

4. Update α from the full conditional,

$$\alpha | \boldsymbol{\pi} \sim \text{Gamma} (a_\alpha + K - 1, b_\alpha - \log \pi_K). \quad (5)$$

5. For $i = 1, \dots, n$, update z_i from the full conditional,

$$z_i | \mathbf{y}_i, \boldsymbol{\pi}, \boldsymbol{\mu}, \Sigma \sim \text{Multinomial}(\pi_{i1}^*, \dots, \pi_{iK}^*), \quad (6)$$

where $\pi_{ik}^* = \pi_k N(\mathbf{y}_i | \boldsymbol{\mu}_k, \Sigma_k) / \left\{ \sum_{g=1}^K \pi_g N(\mathbf{y}_i | \boldsymbol{\mu}_g, \Sigma_g) \right\}$.

We repeat these steps for many iterations until apparent convergence of the chain, discarding iterations before apparent convergence. Convergence can be assessed by examining trace plots of the parameters, especially α which is not subject to label switching (Gelman et al., 2013).

3 Results for the Plastics Products Industry from the CMF

In Section 5 in the main text, we presented results for the concrete industry. Here, we present results of similar analyses for the plastics products manufacturing industry. The qualitative conclusions about the methodology are essentially identical for both industries.

We consider three scenarios when setting $\boldsymbol{\pi}^*$. The first scenario assumes MAR, i.e., $\boldsymbol{\pi}^* = \boldsymbol{\pi}^m$. The second scenario assumes NMAR with nonrespondents that tend to be small establishments. Thus, we inflate π_k^m for components with small values of Δ_k . The third scenario assumes NMAR with nonrespondents that tend to come from the clusters with high

values of the variables. Thus, we inflate π_k^m for components with large values of Δ_k . For all scenarios, we use the mixture normal model with $\Sigma_k = \sigma_p \mathbf{I}_p$ for all k , where $\sigma = 0.3$. The MAP iteration has eight non-empty components.

Figure 1 displays pairwise scatter plots for the plastics products industry after imputation using MAR. Most of the density is concentrated in fewer than half of the components. As desired in the MAR scenario, the imputed values for nonresponse have essentially the same distribution as \mathbf{Y}_t , which suggests that the mixture normal model describes the distribution of \mathbf{Y}_t reasonably well.

Figure 2 and Figure ?? display results for the plastics products industry after imputation for the two NMAR scenarios. In Figure 2, the imputed values for nonrespondents are primarily in regions with small values of \mathbf{Y}_t . This scenario is obtained by setting π_k^m to be large for $k \in (5, \dots, 8)$ (the low-ranking components), taking mass away from the other components. Similarly, in Figure 3 we set π_k^* to be large for the top-ranking components, generating imputed values for nonrespondents that tend to be in the region of high values of \mathbf{Y}_t .

Figure 4 displays $(\theta_\delta^{(s)}, \tau_\delta^{(s)}, \rho_\delta^{(s)})$ for the three missingness scenarios in the plastics products industry data as a function of δ . In the MAR scenario, we produce greater accuracy when generating the $R = 5$ imputed datasets from the model estimated with $(\mathbf{Y}_t, \mathbf{Y}_{f,\delta}^{(s,j)})$ than the model estimated with \mathbf{Y}_t . This is expected, since the former model is based on more observations than the latter, and both \mathbf{Y}_t and $\mathbf{Y}_{f,\delta}^{(s,j)}$ are representative of the nonrespondents under MAR. In contrast and also as expected, in the NMAR scenarios the values of the utility measures are much smaller when the imputation models are estimated only with $\mathbf{Y}_{f,\delta}^{(s,j)}$, which is representative of the nonrespondents whereas \mathbf{Y}_t is not. The values of the utility measures are generally smaller for the MAR scenario than for the NMAR scenario.

In the MAR case, a 25% nonresponse followup sample decreases $\theta_\delta^{(s)}$ by around 11.6 points (from .300 to .219), and a 50% nonresponse followup sample decreases $\theta_\delta^{(s)}$ by around 10.1 points (from .300 to .199). On the other hand, in the NMAR scenarios, collecting data from 25% or 50% of the nonrespondents results in massive accuracy improvements: $\theta_\delta^{(s)}$ drops from values of at least 1.0 when $n_f = 0$ to values in line with those for the MAR scenario. For all scenarios, beyond $\delta = .50$ the gains seem unlikely to be worth the cost. We reach similar qualitative conclusions for all three utility measures.

4 Acknowledgments

Any opinions and conclusions expressed in this article are those of the authors and do not necessarily represent the views of the U. S. Census Bureau. All results have been reviewed to ensure that no confidential information is disclosed.

References

- Gelman, A., Carlin, J., Stern, H., Dunson, D., Vehtari, A., and Rubin, D. (2013). *Bayesian Data Analysis*. Chapman & Hall/CRC Texts in Statistical Science. Taylor & Francis.
- Ishwaran, H. and James, L. F. (2001). Gibbs sampling methods for stick-breaking priors. *Journal of the American Statistical Association*, 96(453).




Figure 1: Pairwise scatter plots with the results for the plastics products industry from the 2007 CMF, for the MAR imputation scenario. Observed points are plotted as gray hollow circles. The black filled circles on the lower diagonal are the imputed points. The colored circles on the upper diagonal are the 95% quantile ellipses of the fitted clusters, with color intensity proportional to the mixture probabilities. Graph copied from the output of the NIMC R tool. Axes removed to prevent disclosure of information about the magnitude of the data.

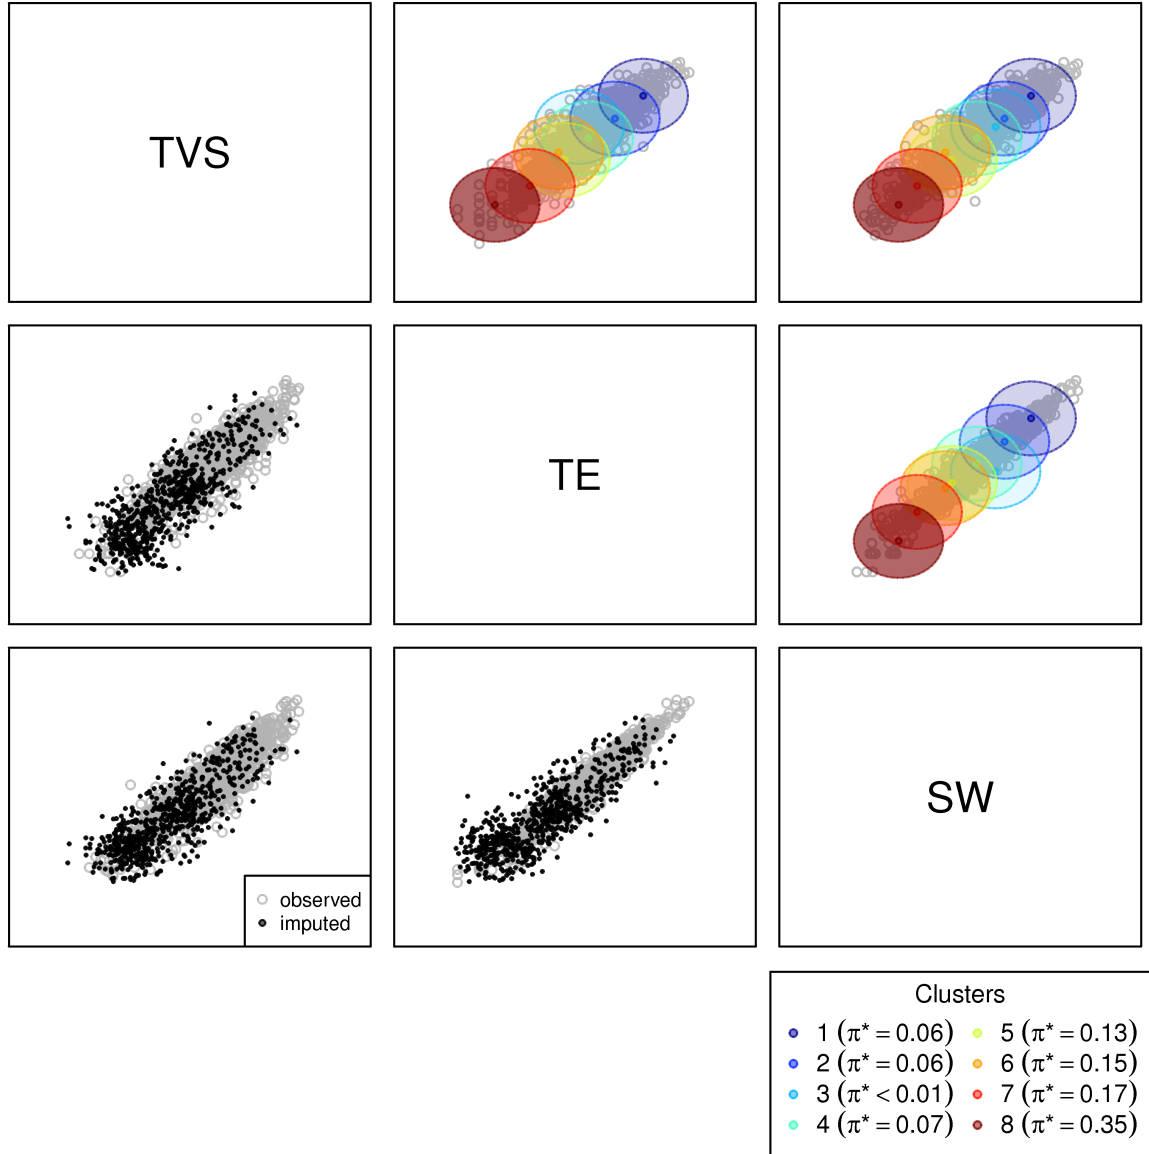


Figure 2: Pairwise scatter plots with the results for the plastics products industry from the 2007 CMF, for the NMAR imputation scenario with higher probabilities for bottom ranked clusters. Observed points are plotted as gray hollow circles. The black filled circles on the lower diagonal are the imputed points. The colored circles on the upper diagonal are the 95% quantile ellipses of the fitted clusters, with color intensity proportional to the mixture probabilities. Graph copied from the output of the NIMC R tool. Axes removed to prevent disclosure of information about the magnitude of the data.




Figure 3: Pairwise scatter plots with the results for the plastics products industry from the 2007 CMF, for the NMAR imputation scenario with higher probabilities for top ranked clusters. Observed points are plotted as gray hollow circles. The black filled circles on the lower diagonal are the imputed points. The colored circles on the upper diagonal are the 95% quantile ellipses of the fitted clusters, with color intensity proportional to the mixture probabilities. Graph copied from the output of the NIMC R tool. Axes removed to prevent disclosure of information about the magnitude of the data.

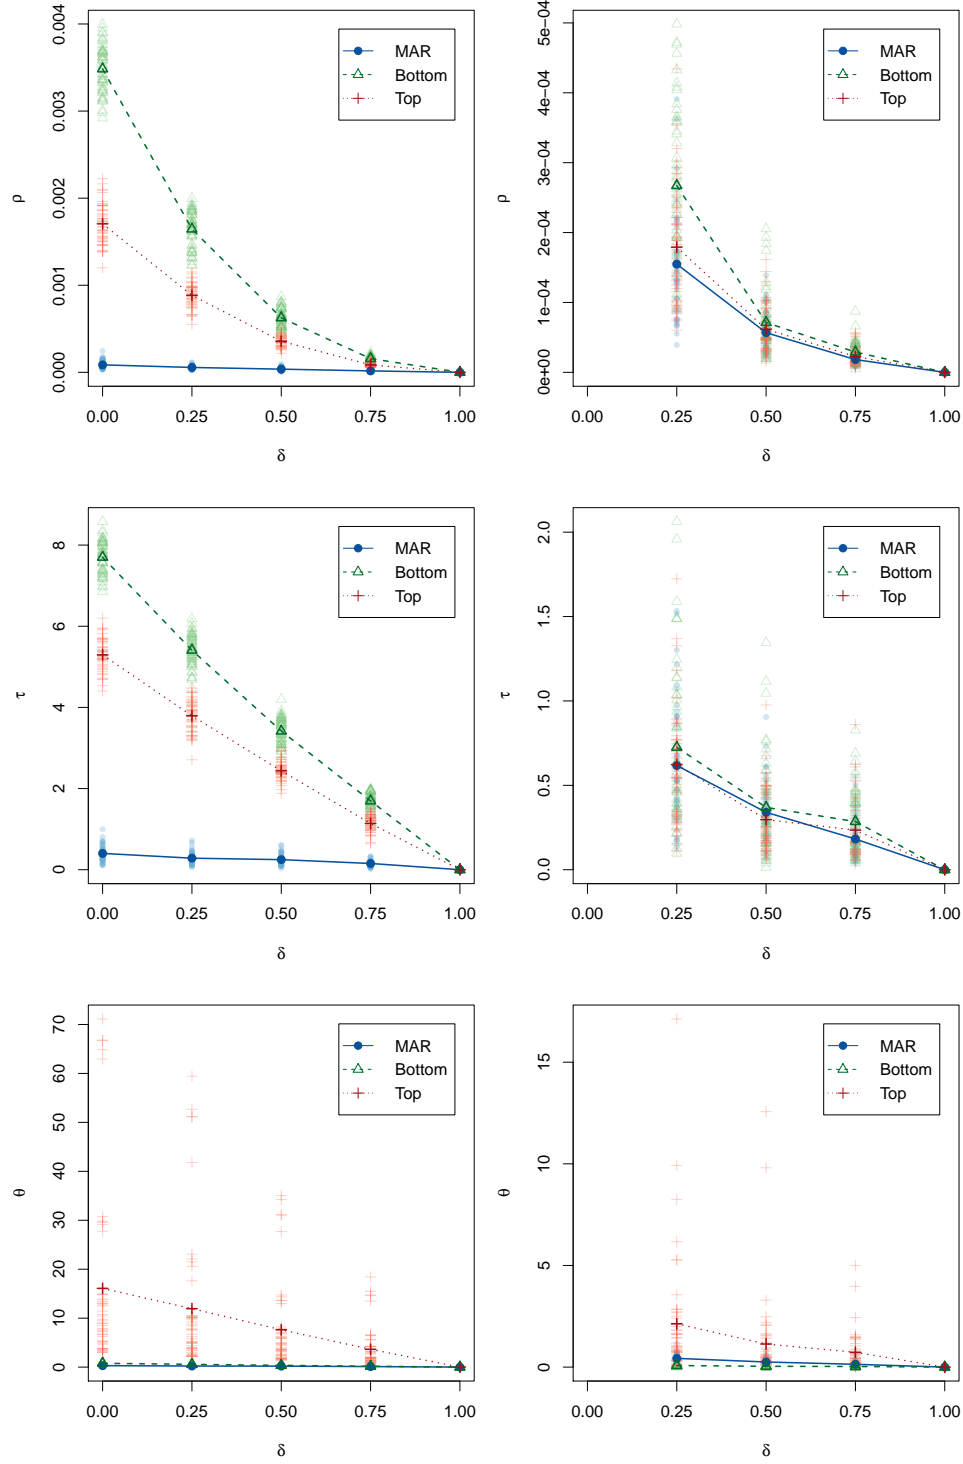


Figure 4: Summary of utility measures for the three scenarios considered for the plastics products industry from the 2007 CMF. The results for ρ , τ and θ are on the first, second and last row, respectively. The plots on the right column contain the results with the model fit to observed data and follow-up sample, while the plots on the right contain the results with the model fit to follow-up sample only. The faded points are the individual values for each multiple imputation.